

A Survey of Systematic Evidence Mapping Practice and the Case for Knowledge Graphs in Environmental Health and Toxicology

Taylor A.M. Wolffe,^{*,†,1} John Vidler,[‡] Crispin Halsall,^{*} Neil Hunt,[†] and Paul Whaley^{*,§}

^{*}Lancaster Environment Centre, Lancaster University, Lancaster LA1 4YQ, UK; [†]Yordas Group, Lancaster University, Lancaster LA1 4YQ, UK; [‡]School of Computing and Communications, Lancaster University, Lancaster LA1 4WA, UK; and [§]Evidence-Based Toxicology Collaboration, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland 21205

¹To whom correspondence should be addressed at E-mail: t.wolffe@lancaster.ac.uk.

ABSTRACT

Systematic evidence mapping offers a robust and transparent methodology for facilitating evidence-based approaches to decision-making in chemicals policy and wider environmental health (EH). Interest in the methodology is growing; however, its application in EH is still novel. To facilitate the production of effective systematic evidence maps for EH use cases, we survey the successful application of evidence mapping in other fields where the methodology is more established. Focusing on issues of “data storage technology,” “data integrity,” “data accessibility,” and “transparency,” we characterize current evidence mapping practice and critically review its potential value for EH contexts. We note that rigid, flat data tables and schema-first approaches dominate current mapping methods and highlight how this practice is ill-suited to the highly connected, heterogeneous, and complex nature of EH data. We propose this challenge is overcome by storing and structuring data as “knowledge graphs.” Knowledge graphs offer a flexible, schemaless, and scalable model for systematically mapping the EH literature. Associated technologies, such as ontologies, are well-suited to the long-term goals of systematic mapping methodology in promoting resource-efficient access to the wider EH evidence base. Several graph storage implementations are readily available, with a variety of proven use cases in other fields. Thus, developing and adapting systematic evidence mapping for EH should utilize these graph-based resources to ensure the production of scalable, interoperable, and robust maps to aid decision-making processes in chemicals policy and wider EH.

Key words: systematic evidence map; knowledge graph; evidence synthesis.

Data relevant to assessing the human and ecological health risks associated with exposure to chemical substances are increasingly available to stakeholders (Barra Caracciolo *et al.*, 2013; Lewis *et al.*, 2016). This trend is owed to a variety of factors, including the advent of the Internet and increasingly sensitive analytical techniques (Lewis *et al.*, 2016), regulatory and economic changes (Lyndon, 1989; Pool and Rusch, 2014), demands for increased

transparency (Ingre-Khans *et al.*, 2016), stricter regulatory data requirements (Commission of the European Communities, 2001; United States Environmental Protection Agency, 2016), reform of regulatory reliance on *in vivo* toxicity testing (ECHA, 2016), and a continually growing chemicals industry. The growing pool of available evidence has significant potential for informing regulatory and risk management decision making.

© The Author(s) 2020. Published by Oxford University Press on behalf of the Society of Toxicology.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

Box 1 Glossary of Terms

Database	An organized and structured collection of information (data) stored electronically within a computer system, which allows data to be accessed, manipulated, and updated.
Systematic evidence map (SEM)	A queryable database of systematically gathered evidence (eg, academic literature and industry reports). SEMs extract and structure data and/or metadata for exploration following a rigorous methodology which aims to minimize bias and maximize transparency.
Coding	The process of assigning controlled vocabulary labels or categories (referred to as “code”) to data, which allows comparisons to be drawn despite the heterogeneity of the underlying dataset. For example, extracted data such as “mouse,” “rat,” and “guinea pig” might all be coded as “rodent” for broad comparison.
Query	A request for data from a database. By requesting data that meets a particular set of conditions, users can query a database for a subset of information of relevance to their specific research interests.
Schema	The organizational plan (“blueprint”) for the structure of a database, detailing the entities stored in the database, the attributes associated with those entities, how those entities are related, what data-types can be stored in the database, etc.
Schemasless	Refers to databases which do not have a fixed and predefined schema.
Schema, on-write	Refers to the application of a schema before data is stored (written) to the database.
Schema, on-read	Refers to the application of a schema after data has been written to the database, at the time the data is accessed (read).
Ontology	A shared and reusable conceptualization of a domain which applies a logically related controlled vocabulary to describe the domain concepts, their properties and relations.

Evidence-based approaches aim to minimize the bias associated with cherry-picking an unrepresentative subset of evidence for consideration in the decision-making process. They advocate for robust, transparent consideration of all relevant, available data and are the core of the evidence-based toxicology movement (Hoffmann and Hartung, 2006; Hoffmann et al., 2017). However, locating, organizing, and evaluating all relevant data is challenging when the quantity of that data is very large and growing exponentially.

Systematic evidence mapping is 1 such evidence-based approach to drawing into consideration all data which are relevant to chemicals policy and risk management workflows (see Wolffe et al., 2019). Systematic evidence maps (SEMs) are queryable databases of systematically gathered research (Box 1). They provide users with the computational access needed to organize, compare, analyze, and explore trends across a broader evidence base (Clapton et al., 2009; James et al., 2016) by:

- Collating data from different sources and storing it in a single location, such that users need only query a single database to satisfy their information requirements;
- Extracting unstructured data and storing it in a structured format, such that data can be programmatically accessed and analyzed;
- Categorizing extracted data using controlled vocabulary code, such that evidence can be broadly and meaningfully compared despite its inherent heterogeneity.

SEMs organize and characterize an evidence base such that it can be explored by a variety of end-users with varied specific research interests. The methodology was developed to address some of the limitations of systematic review and has found application in fields where formulating a single, narrowly focused review question is difficult or uninformative (Haddaway et al., 2016; James et al., 2016; Oliver and Dickson, 2016; Wolffe et al., 2019). Similarly faced with this challenge is chemicals policy and the fields which it encompasses, ie, environmental health (EH) and toxicology. It is difficult to frame a single research question with a scope which is simultaneously narrow enough to elicit the synthesis of a coherent conclusion through systematic review, and also broad enough to address the varied information requirements of chemicals policy workflows. This means that potentially several syntheses over multiple systematic reviews are required to facilitate a single decision-making process in chemicals policy. However, the significant demand for time and resources associated with systematic reviews, and the unmatched resource availability of chemicals policy, necessitates a priority setting, or problem formulation process to ensure the most efficient use of systematic review. Thus, systematic evidence mapping provides a valuable first step in this prioritization process, where the identification of emerging trends across the wider evidence base ensures resources can be targeted most efficiently (see Wolffe et al. [2019] for further discussion of the applications of SEMs in chemicals policy).

These issues are likely to become increasingly pressing as the chemicals policy paradigm shifts toward more evidence-based approaches and methods such as systematic review gain prominence. For example, agencies such as the U.S. EPA (EPA, 2018), EFSA (European Food Safety Authority, 2010), and WHO (Mandrioli et al., 2018; World Health Organization, 2019) have already begun to incorporate systematic review in their chemical risk assessment frameworks. Thus, ensuring that evidence synthesis efforts are targeting the most appropriate issues, and that the data collated for synthesis can be accessed for alternative applications, potentially across agencies, is increasingly important.

Interest in the application of SEM methodology for this context is beginning to emerge in the form of SEM exercises targeting chemicals policy issues (Martin et al., 2018; Pelch et al., 2019), various working groups expanding their evidence synthesis activities to include broader scoping and surveillance exercises (NTP-OHAT, 2019; Pelch et al., 2019; The Endocrine Disruption Exchange, 2019; Walker et al., 2018), and conference sessions discussing the potential benefits of SEMs for EH (Beverly, 2019). This emerging interest in SEM methodology, and its ability to facilitate evidence-based approaches, necessitates study of the factors key to its successful adaptation to EH contexts.

Therefore, we seek to understand how SEM databases are built and presented to end-users in fields where the practice is more mature. We hope that contextualizing this understanding within the needs of chemicals policy, risk management, and wider EH research will expedite the development of effective evidence mapping methods in this domain.

To achieve this, we examine the current state-of-the-art and common practices associated with constructing and presenting a SEM database in environmental management, a field with a strong history of systematic mapping publications and method development (Collaboration for Environmental Evidence, 2019c; Haddaway et al., 2016, 2018a; James et al., 2016). We discuss the implications of current practices for EH and highlight the challenges associated with using rigid data structures for storing the highly connected and heterogeneous data associated with the field. We outline the need for more flexible data structures in

Table 1. The Concepts Used to Guide Data Extraction and Subsequent Assessment and Discussion of the Outputs of CEE Systematic Mapping Exercises

Concept	Definition	Metadata Extracted
Data storage technology	How data extracted and collated during the systematic mapping exercise were stored for future exploration	Format in which the systematic map database is presented to users (eg, spreadsheet, relational database, in-text data table, and in-text figure).
Data integrity	How accurately the systematic map is able to represent the raw study data on which it is based	How the relationships between entities (or study attributes) which underpin the raw data are maintained in the systematic map.
Data accessibility	How easy it is for end-users to access the data relevant to their research interests, or the ability of the systematic map to return data relevant to an end-user's queries	The querying mechanisms recommended in the systematic map's study report (eg, filtering table columns and navigating interactive dashboards).
Transparency	The ability of end-users to verify how the systematic map represents the raw study data on which it is based, ie, whether the map maintains a link between raw extracted data and eg, controlled vocabulary code.	Whether the map maintains a link between raw extracted data and controlled vocabulary code (eg, map presents code-only, map presents raw data and code), and how this link is maintained.

EH SEMs and introduce the concept of “knowledge graphs” as an effective and intuitive model for the storage and querying of highly connected EH data. Finally, we discuss graph-based SEMs in the context of current, complementary efforts in the development of toxicological ontologies, outlining the future of systematic evidence mapping for regulatory decision making.

MATERIALS AND METHODS

Survey of published Collaboration for Environmental Evidence SEMs. We identified a dataset of exemplar SEMs for analysis: the complete set of SEMs of the Collaboration for Environmental Evidence (CEE). These maps were chosen because of CEE's role in pioneering the adaptation of systematic mapping methodology from the social sciences (Clapton et al., 2009; James et al., 2016). Through example (Collaboration for Environmental Evidence, 2019b), communication (Collaboration for Environmental Evidence, 2019a), published guidance (James et al., 2016), and reporting standards (Haddaway et al., 2018b), CEE advocate for systematic mapping and represent an ongoing case study for how the methodology can be developed as a policy and decision-making tool. Understanding how systematic map outputs serve this function, and what methodological adaptation is required to produce these outputs, is vital for successfully applying the methodology in EH. Thus, the outputs (ie, the queryable databases) of CEE's more firmly established systematic mapping practice were surveyed.

All CEE systematic maps completed before July 2019 were identified in the CEE Library (<http://www.environmentalevidence.org/completed-reviews>, last accessed July 2019). The study reports and the Supplementary information for these maps were downloaded and key metadata extracted, including title, authors, publication date, and map objectives (Supplementary Table 1). Metadata regarding the output of the systematic mapping exercises were then gathered and assessed in duplicate by T.A.M.W. and P.W. using a data extraction sheet which asked open-ended questions relating to 4 key themes of analysis: data storage technology; data integrity; data accessibility; and transparency (Table 1). These themes were developed in discussion among J.V., T.A.M.W., and P.W.

“Data storage technology” concerns the software used to construct the systematic map databases and their associated data storage formats.

“Data integrity” concerns the structures of the CEE maps. Although an important aspect of data integrity, appraising the

data extraction efforts of mappers (ie, confirming that the data extracted, coded, and stored in the database are an accurate representation of their raw counterparts in the primary literature) was beyond the scope of this exercise. Rather than verifying the data, *how* that data are represented (regardless of *what* is represented) by the systematic map database output was assessed by focusing on the ability of the systematic map to maintain the relationships which underpin these data. For example, a mapper may have extracted data from a study which investigates outcomes in a population. Although the mapper may have extracted data such as “outcome x” and “population y”—the manner in which the database structures and organizes these data will determine whether end-users can decipher that “outcome x” is somehow related to “population y.”

“Data accessibility” concerns the capacity for CEE's systematic maps to facilitate data exploration by end-users. Systematic maps are research products in their own right (Haddaway et al., 2016). They should therefore present end-users with a means of programmatically accessing and querying the data they store, such that trends in potentially large datasets can be quickly identified with minimal manual effort. Accessibility is an important consideration when producing maps for an audience of varied technical skill, where ensuring that the map is accessible for nonspecialist users should not compromise the ability of more technical users to run complex queries. Therefore, the extent to which CEE systematic mapping exercises consider accessibility from the perspective of users was surveyed by extracting eg, details on the level of guidance provided to end-users wishing to query the systematic map database, and recording P.W. and T.A.M.W.'s experience of interacting with and querying the maps.

Finally, “transparency” concerns how systematic maps facilitated an end-user's ability to validate the extent to which the data presented in a map represents the data in the primary research. This was achieved by determining whether the map preserved a link between raw data and assigned controlled vocabulary labels/categories (“code” - see Box 1).

T.A.M.W. and P.W. independently noted answers to the data extraction questions before discussing and agreeing on an aggregate, consensus view. This was to contribute to comprehensive coverage of potential discussion points in relation to each theme. These aggregate assessments are presented in Supplementary Tables 1–6 and are used to evidence the state-of-the-art in terms of producing queryable systematic map databases for exploration of the environmental management




Figure 1. Publication history of CEE SEMs indicating the number of maps published per year. The year in which the CEE guidance on systematic mapping methods was published (2016) is marked on the corresponding bar (James et al., 2016).

literature. Their contents are referenced throughout the Results and Discussion sections of this survey.

RESULTS

Twenty-one systematic maps covering a variety of topics were identified in the CEE library, published between October 2011 and January 2019 (Figure 1).

The aggregated, narrative assessments of each CEE systematic map can be found in [Supplementary Tables 1–6](#). The extracted data and aggregated assessments for each CEE systematic map are organized as follows:

- [Supplementary Table 1](#)—Bibliographic information
- [Supplementary Table 2](#)—Data storage technology
- [Supplementary Table 3](#)—Data integrity
- [Supplementary Table 4](#)—Data accessibility
- [Supplementary Table 5](#)—Transparency
- [Supplementary Table 6](#)—Additional notes

Excluded Maps

Two systematic maps (Johnson et al., 2011; Mcintosh et al., 2018) are assessed in the [Supplementary information](#) but are excluded from further analysis, as neither provided a database output which could be analyzed using our framework. Mcintosh et al. (2018) yielded a null result and therefore provided no database; Johnson et al. (2011) predated CEE’s Environmental Evidence journal and its definition of systematic mapping and, although it is included in the CEE library, presented only in-text tables without an accompanying database.

Data Storage Technology

Two different data storage technologies are used in the outputs of CEE systematic mapping projects: spreadsheets constructed in Microsoft Excel ($n = 14$); and relational databases constructed with the Microsoft Access relational database management system ($n = 5$). One mapping exercise used both of these technologies to present its outputs in 2 different formats (Haddaway et al., 2014). The 2 versions of Haddaway et al. (2014) appear to be identical except that the spreadsheet version includes the results of a critical appraisal process where the relational database version does not. As the spreadsheet version presents the more complete dataset, Haddaway et al. (2014) has been coded as a spreadsheet-based systematic map for the purposes of this survey (see [Supplementary Table 2](#), discussed in the “Data

Integrity” section). A brief description of each identified storage technology can be found in [Table 2](#).

Data Integrity

A single, flat data table (2-dimensional array of rows and columns) was the output for the majority (84%) of CEE systematic maps (16 of 19 maps surveyed, ignoring any look-up tables housing controlled vocabulary code). 80% (4 out of 5) of the maps using the relational database storage technology were also structured as a single, flat data table.

Three maps presented more than 1 table. Two presented at least 2 tables in separate files which were not formally related to each other (Haddaway et al., 2018a; Sola et al., 2017), and 1 presented multiple tables which were related to each other in a 1:1 manner within a relational database. Systematic maps were considered to be stored in more than 1 table if there was limited overlap of the data fields housed in each table ie, if querying the map required accessing information from more than 1 table. Sola et al. (2017) is an example of this, providing the results of its quality appraisal process separately to the data it extracted and coded from the literature—thus any queries investigating critical appraisal in conjunction with another variable require the user to access information from both tables. This distinction was required because some maps, Haddaway et al. (2014) and Randall et al. (2015), presented their outputs in multiple tables, but the additional tables were simply subsets of the most complete table (ie, there was no data in the smaller tables not already present in the largest table).

Several studies included in the systematic maps contained multiple potential values for a particular attribute eg, if a single study had multiple populations and/or multiple outcomes.

Common strategies for maintaining relationships between such data in the tables of CEE maps included “expanding rows” ($n = 6$), “expanding columns” ($n = 2$), or a combination of both ($n = 5$) (see [Figure 2](#)). The remaining 6 maps either did not present/extract studies with multiple potential values per attribute ($n = 1$) or opted to house multiple values within a single cell of the table ($n = 5$, discussed further below).

“Expanding rows” refers to the practice of structuring a data table in *long form*: recording an entity over multiple rows. In long-form tables, a study investigating eg, 3 different outcomes might be recorded over 3 different rows. Although the data entered under the “outcome” data field might be unique in each of these 3 rows, the data for all other attributes will be repeated ([Figure 3A](#)).

In contrast, “expanding columns” describes the practice of structuring a data table in wide form; expanding what would be considered a single data field in long-form tables across several columns. Thus, all unique values associated with the data field can be recorded across a single row, eg, a study reporting 3 different outcomes might be recorded across a single row if the “outcome” attribute is split into 3 unique columns (eg, “outcome 1,” “outcome 2,” and “outcome 3”) ([Figure 3B](#)).

The other strategy for presenting related data in a table was to record multiple values within a single cell for multiple data fields ($n = 11$), whereas 1 map presented multiple values per cell for only a single data field within the database (this distinction matters for reasons we discuss below). The practice of presenting multiple values in a single cell of the database was observed for most (5 of 6) of the maps which avoided expanding row/column structure, and similarly for most (5 of 6) of the maps adopting a long form, expanded row structure.

Table 2. Description of the Storage Technologies Used by CEE Systematic Maps

Storage Technology	Description
Spreadsheets	Spreadsheets are stand-alone applications which offer functionality for end-users wishing to explore and/or manipulate data (Zynda, 2013). A spreadsheet stores data in the cells of 2-dimensional arrays made up of rows and columns. By referencing the coordinates of cells in mathematical formulae, spreadsheet applications such as Microsoft Excel facilitate analysis, transformation, and visualization of tabular data. Although designed and optimized for quantitative data and accounting applications, spreadsheets are commonly used for storing and organizing data in a variety of research contexts, including systematic mapping exercises.
Relational databases	A relational database uses several formally described tables to organize data. Each table stores instances of an entity (across rows), described by a series of attributes (columns). In contrast to storing data in a single, flat data table, relational databases are able to preserve the connection between related entities. These connections are predefined and created through a system of referencing unique identifiers (primary/foreign keys) in corresponding tables. This allows users to enrich their queries with connected information, such that more complex questions can be asked of the evidence base (Elmasri and Navathe, 2013).




Figure 2. The number of CEE systematic maps that are structured with expanding rows and/or expanding columns as a means of preserving data relationships. Maps using the relational database storage technology are presented on the left, while maps using the spreadsheet storage technology are presented on the right. In addition, the numbers of systematic maps which store multiple values within a single cell of their data table/s are indicated by solid shading, whereas those that do not are indicated by patterned shading.




Figure 3. Illustrative example of how “expanding-rows” (A) and “expanding columns” (B) are used in long-form (A) and wide-form (B) tabular data structures, respectively.

Data Accessibility

Eighteen of 19 surveyed systematic maps presented users with static data visualizations within their study reports (eg, bar charts, tables, and heat maps) as a means of accessing trends within the evidence. Six systematic maps additionally provided users with an open-access interactive data visualization dashboard, such that users could choose trends for exploration within the map. Four of the 6 maps supplied comprehensive guidance and/or instruction for users wishing to interact with the visualization dashboard.

Far fewer mapping exercises provided any such comprehensive guidance for querying their database output, with only 2 of 19 maps providing a detailed help file for users wishing to query the database (Haddaway et al., 2014; Randall and James, 2012). This was also seen in mapping exercises presenting guidance on interacting with their data visualization dashboards, none of which provided equivalent detailed guidance for querying the underlying database. Instead, 6 CEE systematic maps dedicated only brief discussion to querying within the text of their study reports, leaving 11 maps which offered no discernible guidance.

Where provided, the querying practices identified in user guidance/instruction were “filtering,” “sorting”/“ordering,” “searching,” or some combination thereof (see Supplementary Table 4). Specific examples of queries which could be run against the database were rarely provided in such guidance, with only 2 of 19 maps providing an illustrative example of how a user’s plain-text question is translated into querying the database (Haddaway et al., 2014; Randall and James, 2012), and a further 2 of 19 making only brief mention of how a specific data field might be filtered (Cresswell et al., 2018; Randall et al., 2015). None of the maps reported the queries or querying processes used to generate visualizations or analyses. Two maps (Cheng et al., 2019; McKinnon et al., 2016) indicated that an additional data processing step had been conducted eg, using the statistical programming language R. Cheng et al. (2019) provided a link to the code used for this analysis, however the link was broken at the time this survey was conducted.

Transparency

Thirteen of 19 surveyed CEE systematic maps presented only the controlled vocabulary code which was used to classify the data of interest, not recording the raw data itself in the map. Six of 19 maps maintained a link between this code and the raw data/the coders’ interpretation of the raw data. Approaches to this included using data fields which contain free-form text alongside the controlled vocabulary terms applied to categorize this free text (5 of 6 maps, Macura et al., 2015), and providing the location of the raw data within the original study report represented as code in the systematic map (1 of 6 maps, Haddaway et al., 2015).

Seventeen of 19 CEE mapping exercises provided a codebook. Codebooks were generally supplied separate to the systematic map database, in a different file and/or format ($n = 14$), although some incorporated codebooks into the database as either lookup tables ($n = 1$, Leisher et al., 2016), or separate spreadsheets within the same workbook as the systematic map ($n = 2$, Bernes et al., 2015, 2017).

Codebooks largely presented the controlled vocabulary terms used to code study attributes (12 of 17) but did not always provide this detail (5 of 17). For codebooks which did provide controlled vocabulary terms, a narrative description or discussion of the potential types of data which might be assigned certain codes was presented in only 2 of the codebooks.

Relationships between controlled vocabulary terms were generally omitted from codebooks and/or the systematic map databases themselves, except for 1 map which structured its code as a hierarchy of nested terms (Haddaway et al., 2015).

DISCUSSION

CEE has been a driving force for the introduction of systematic mapping to the environmental sciences. Their maps act as case studies for adapting evidence-based methodologies to other fields. CEE’s involvement of stakeholders in their systematic mapping approach has undoubtedly resulted in outputs of value to those stakeholders and their specific research contexts (Haddaway and Crowe, 2018). The following discussion does not critique the use of CEE’s systematic maps for their intended purposes, but instead takes the perspective of EH applications to identify transferable aspects of current practice and remaining challenges.

Systematically Mapping the EH Evidence Base: General Considerations

EH data are complex, heterogeneous, and highly interconnected (Vinken et al., 2014). Chemical risk assessment and risk management seek to understand the outcomes which result from these complex connections—synthesizing evidence of varied resolution and origin eg, considering in combination evidence from bio- and/or environmental monitoring, *in vitro*, *in vivo*, *in silico*, and/or epidemiological studies (Martin et al., 2018; Rhomberg et al., 2013; Vandenberg et al., 2016).

The relationships which hold the disparate EH evidence base together are vital for building a more complete understanding of toxicity. These relationships underpin adverse outcome pathways (ie, how molecular initiating events lead to apical outcomes through a causal pathway of connected key events [Edwards et al., 2015]), quantitative structure-activity models (ie, how the chemical structure of a substance can be quantitatively related to its physicochemical properties and biological activity [Schultz et al., 2003]), read-across applications (ie, where predictions for data-poor substances are based on structurally related data-rich substances) and other key components of chemicals policy workflows. Such relationships are also vital for understanding the impacts of real-world exposures to mixtures of chemical substances (Sexton and Hattis, 2007).

The interconnectedness of the EH evidence base means that even if SEM methodology is used to explore just a subset of EH research, or to facilitate just 1 component of chemicals policy workflows—the data collated, extracted, and coded are likely to be of relevance to a myriad of alternative EH research interests and chemicals policy applications. Thus producing “multi-purpose,” interoperable EH SEMs that can be queried according to a variety of specific use cases is the most resource-efficient means of implementing the methodology.

However, many of the complex relationships constituting the EH evidence base are unknown to individual users, who will only have cognitive access to part of the total knowledge space in a given domain. Thus, in addition to facilitating the identification of trends which are based on relationships already known to users, EH SEMs should also facilitate the identification of relationships which are unknown to users. This would enable a more highly resolved and customizable querying process which extends beyond the user’s personal understanding of the domain, adding valuable connected contextual information with which to explore and interpret trends. It is this value, gained through accessing as well as exploring relationships—




Figure 4. A, The relationship between attribute A and entity 3 is explicit in the formal structure of the array. However, the relationship between attribute A and attribute C is implicit and has to be inferred by the user from features external to the table eg, conventions around interpreting tabular data. The external conventions are not part of, or known to, the table and may not be known to the user. B, For example, a user may (in this case, correctly) infer that “sex” is a property of “species” and not “outcome,” but this inference is made using external conventions and contextual understanding—the relationship is not in fact known to the table. All the table can assert is that each entity 1 through 6 has a relationship to properties of sex, age, species, and outcome, respectively.

along with the inherent complexity of those relationships—which makes the flat and rigid tabular data structures currently characterizing CEE systematic maps ill-suited to the task of systematically mapping EH data.

Limitations of Current Evidence Mapping Practice: Data Storage and Structure

Data storage is the fundamental component required for creating a systematic map database, underpinning many of the themes assessed in this survey. This discussion focuses on issues of data storage technology and its close relationship with data integrity.

Use of spreadsheets (and other flat data tables). The majority of CEE systematic maps are stored and structured as flat data tables, mostly as spreadsheets. Tables are a simple, familiar, and robust means of structuring data. However, maintaining relationships within a 2-dimensional array of rows and columns can be challenging. This is because the only explicit relationships in a 2-dimensional array (single table), are between the attributes (columns) and the entities (rows). Any relationships which exist between columns/attributes in a table can only be inferred by the user (Figure 4). We found making such inferences a challenge when surveying systematic maps of research outside of our own fields of expertise (see Supplementary Table 3). The prior knowledge required to successfully navigate data relationships within tabular maps limits their accessibility for less specialized users.

A variety of techniques were employed by CEE maps for maintaining the relationships between attributes, and for

allowing attributes to record multiple values. Of particular note were the practices of expanding columns to produce wide-form tables, and of housing multiple values within a single cell. Although expanding columns and/or housing multiple data entries in single cells do not threaten data integrity when applied to only 1 single attribute (see Thorn et al., 2016, Supplementary Table 3), a loss of referential integrity was noted for maps implementing this practice for multiple attributes.

Such loss is illustrated in Figure 5, whereby column expansion (Figure 5A), and similarly multivalued cells (Figure 5B), falsely assert data relationships unrepresentative of the raw extracted data. Loss of referential integrity is acknowledged by Neaves et al. (2015), where the authors highlight falsely asserted interattribute relationships as a limitation of their mapping exercise.

The alternative strategy used by CEE systematic mappers when structuring data as a flat table was row expansion. Although advantageous for maintaining referential integrity, these long-form data structures can be challenging to process. They can create confusion for end-users interpreting what the study “unit” (entity) which constitutes a new row in the data table is (see Supplementary Table 3). Users must also be cautious of duplicates when querying specific data fields within the table. Duplicating data can also increase the risk of data-entry errors for systematic mappers tasked with manually populating a long-form table, resulting in inconsistencies.

In summary, the spreadsheet storage technology is an unsuitable long-term solution for EH SEMs, with wide-form tables potentially compromising data integrity, and long-form tables being impractical and/or error-prone.




Figure 5. A, Loss of referential integrity resulting from the column expansion of more than 1 study attribute (data field). The recording of multiple populations and multiple outcomes on a single row compromises the ability of users to decipher which population was affected by which outcome. The table asserts that both populations (mice and rats) were affected by all 3 outcomes (reduced birth weight, tumors, and behavioral changes), respectively—which may not be truly representative of the raw data, compromising data integrity. B, This is similarly observed when multivalued cells are used for more than 1 study attribute.

Use of relational databases. Many of the discussed challenges associated with implementing systematic maps as flat data tables or spreadsheets are addressed by relational databases—the alternative storage technology identified in current systematic mapping practice (see Table 2). Relational databases divide entities into their own, referenceable tables—allowing links between related entities to be created and maintained. These links are coded into the database itself, and therefore do not rely on an end-user’s implicit understanding of external conventions to correctly interpret.

The structure of a relational database is organized in an *on-write* schema, which is effectively a “blueprint” for the database (Karp, 1996); ie, the schema defines what constitutes an entity and therefore a data table, which attributes describe an entity, how an entity is related to other entities and therefore how data tables must reference others, all before data are stored. This necessitates a sound understanding of both the data to be stored in the database, and also the potential applications of the database. In fact, the optimization of end-users’ capacity to query the database for a particular application is a key driver of schema design (Blaha et al., 1988).

The “schema first, data later” (Liu and Gawlick, 2015) approach of relational databases requires a more detailed level of prior knowledge regarding the structure of the evidence and/or the applications of the database. This is problematic for EH SEMs for several reasons.

First, the potential applications of an EH SEM are varied. Even where a specific use case is known, an EH SEM should at least avoid restricting access to the evidence base for alternative uses. Second, SEM methodology advises against making decisions which are based on post hoc assessment of included

studies (James et al., 2016). However, without this assessment it is difficult to design a schema capable of housing all the entities and relationships likely to arise from the varied study designs and/or evidence streams collated through an EH SEM exercise. Even if this prior assessment were advocated by SEM methodology and did not lead to the introduction of bias or inconsistencies, there would likely be far too much data for mappers to feasibly consider in the design of an EH SEM’s schema.

Third, SEMs are currently constructed by human mappers, who screen, assess, and extract data from 1 included study at a time. In this manner, mappers’ understandings of the relationships between entities are limited to the level of the individual study. Thus, it can be difficult to design a schema able to appropriately account for relationships which occur at an interstudy level, compromising end-users’ ability to query these relationships. For example, a one-to-many relationship between population and outcome entities may be appropriate at the level of the individual study, where a single population can be investigated for many outcomes. However, at the evidence-base level, a particular outcome may in fact have been reported by many studies, and therefore investigated in many different populations—making a many-to-many relationship between population and outcome, and a schema capable of representing this relationship, more appropriate. Alternatively consider the relationships between adverse outcomes along a causal pathway. Although a relationship between eg, Outcome A and Outcome C might become apparent at the evidence base level, mappers may only have access to relationships between eg, Outcome A and Outcome B, or Outcome B and Outcome C—which occur at the individual study level.




Figure 6. (A) Knowledge captured in unstructured, textual formats e.g. scientific articles, is distributed and programmatically inaccessible. (B) This knowledge can be structured in an intuitive and machine-readable way as a series of semantic subject-predicate-object triples – where entities are the subjects and/or objects and the relationships between entities are the predicates. (C) Entities can be stored as the nodes of a graph. The semantic value of the relationships between entities are preserved and stored as edges. The graph can continue to grow to produce a queryable representation of all knowledge on a topic (see Figure 7).

Finally, the growing volume and scope of EH data means that even if it were possible to devise a schema capable of accounting for all study designs that exist at present, new, and emerging study designs would soon out-date the schema, necessitating laborious, and potentially error-prone schema migration (Segaran et al., 2009).

Avoiding these issues and attempting to balance the rigidity of a schema with the fluidity or heterogeneity of the data it organizes forces mappers to implement work-arounds (eg, compromising the resolution of SEMs), the likes of which might compromise the utility of SEMs for chemicals policy applications (see Supplementary File 1).

Overcoming the Limitations of Spreadsheets and RDBs: Knowledge Graphs for Mapping EH Evidence

Expanding and enriching the application of SEMs to varied EH research problems requires moving away from the rigidity of tabular data structures and their predefined relationships. Instead, SEMs in EH should utilize more flexible, *schemaless* data models and storage technologies. We believe this flexibility is offered by knowledge graphs and associated graph-based data storage technologies.

Knowledge graphs. The scientific knowledge codified in a study report can be readily formalized as a set of subject-predicate-object “triples.” These triples can be stored as mathematical “graphs” (nodes and edges) where the nodes are the entities (subjects and objects) and the edges are the predicates, or relationships, between

the subjects and the objects (see Figure 6). Because the graph is a direct representation of the semantic content of the studies being stored, it can be said to represent the knowledge captured in the study—hence “knowledge graph” (Ontotext, 2019b).

In graph database implementations, data are stored as nodes and relationships are stored as edges. Unlike the relational model, the graph model regards relationships as first-class entities, and keeps them alongside the values they connect. Rather than “artificially” creating relationships through cross referencing primary and foreign keys in data tables, graph databases natively store relationships, preserving their semantic value, and making them accessible to queries (Figure 6 and 7) (Robinson et al., 2015). This is particularly valuable when the relationships underpinning data cannot be directly characterized *a priori*, or when the relationship between 2 pieces of information (nodes) can only be inferred through traversal of relationships which indirectly connect those nodes (Ontotext, 2019c) (eg, the inferred causal relationship between “Chemical X” and “Tumours” in Figure 7).

The graph model’s flexibility and emphasis on relationships allows it to accommodate new developments in EH research. Data produced by studies of novel design can be incorporated among, and related to, preexisting data in the database without needing to update schema and subsequently migrate data (Robinson et al., 2015). This is illustrated in Figure 7 which expands the amount of data populating the graph in Figure 6.

Knowledge graphs are already being exploited in other fields centered around the analysis of highly connected data (Ghrab




Figure 7. Storing relationships as first class entities allows knowledge graphs to continue to grow and expand without needing to revise schema and migrate data. This flexibility is particularly useful when relationships between entities cannot be characterised *a priori*.

et al., 2016). Notable use cases for graphs include: mapping complex networks of biological interaction (Aggarwal and Wang, 2010; Have and Jensen, 2013; Pavlopoulos *et al.*, 2011); representing chemical structures (Aggarwal and Wang, 2010); tracking communication and transaction chains for fraud detection (Castellort and Laurent, 2016; Sadowski and Rathle, 2015); feeding recommendation engines for online retailers (Webber, 2018); facilitating highly customized outputs for social media platforms (Gupta *et al.*, 2013; Weaver and Tarjan, 2013); promoting a more proactive service from search engines (Singhal, 2012); and many more. The key commonality between these applications is the identification of trends or patterns of information that facilitate the generation of new knowledge that is actionable or of value to decision-making.

Schemaless data storage and data exploration. As relationships are stored as queryable, first-class entities—the schema which implicitly structures data begins to emerge naturally and can be discovered and exploited by knowledge finding applications *on-read* (Janković *et al.*, 2018; Kleppmann, 2017).

In CEE’s current systematic mapping practice, trend exploration is predominantly reliant on filtering columns of a data table for specific values of interest. This requires that users are familiar with the structure of the database *ie*, they know which columns house values of interest, what those values of interest are, and that their interests align with the data model imposed by the tabular map. By comparison, graphs are amenable to some ambiguity in a user’s query. Beyond the potential existence of an entity of interest, users do not require prior knowledge of the graph’s structure, or the relationships connecting the entity of interest to others, to successfully gain an understanding of the graph space around that entity. This facilitates the building of data models which contextualize this understanding within a particular application.

In current systematic mapping practice, data models are closely tied to the data storage mechanism and its structure. Knowledge graphs do not fix data models *on-write*, separating data models from data storage—thus it is possible to apply multiple models to the same graph, optimizing access to the evidence base for a variety of interests and queries. Changes can also be readily incorporated into these data models without migrating the underlying data they access.

Ontologies. A key component of wider data modeling activities is the development of domain-specific ontologies. An ontology is an agreed upon and shared “conceptualization” of a domain (Dillon *et al.*, 2008), comprising a formal specification of terms used for describing knowledge and concepts within a domain and their relationships to each other, expressed through a standardized controlled vocabulary (Ashburner *et al.*, 2000; National Center for Biomedical Ontology, 2019). Developing domain-specific ontologies closely mirrors the coding step of systematic evidence mapping, which is designed to conceptualize the evidence base through organizing extracted data using a controlled vocabulary of terms.

In knowledge graph applications, ontologies are stored as data themselves (Noy and Klein, 2004)—forming an additional “layer” within the graph. Raw extracted data stored in the graph can be viewed as instances of an ontology’s classes. By using data models to bind nodes of raw data to the nodes of a suitable ontology, users can navigate the evidence base through this ontology—but do not lose the ability to access the underlying raw data relevant to more highly resolved queries. Furthermore, maintaining a link between raw data and the controlled vocabulary code of a shared toxicological ontology serves to promote transparency, interoperability (Hardy *et al.*, 2012), and the development of training sets for machine-learning classifiers.

However, these concepts are underexplored in current evidence mapping practice where the majority of maps presented code in lieu of raw extracted data. This compromises transparency and limits users' ability to query data at variable resolution. In addition, coding vocabularies were rarely descriptive of the relationships that linked 1 term to another, with only 1 map organizing its code as a hierarchy of nested terms (Haddaway *et al.*, 2015). Where relationships between code were implied, this was generally stored in separate codebooks (ie, not as data within the database)—requiring users to consult a separate document for interpretation.

Other Lessons From Current Systematic Evidence Mapping Practice

Studying the key features of a systematic map database, ie, storage technology and the data structuring choices available for those technologies, highlights the need to pursue more flexible, schemaless approaches when adapting the methodology for EH. We have identified knowledge graphs as the technology capable of providing this flexibility. Although briefly covered in the above discussion, this survey identified additional aspects of current evidence mapping practice which are worthy of discussion.

Data accessibility, user-interfaces, and map documentation. A queryable database is the main, but not sole, output of mapping exercises. All CEE maps are accompanied by a study report which details methodology, presents key trends through data visualization, and/or describes further research needs. These accompanying reports can be thought of as documentation for their database products. In the context of software development, documentation is a formal written account of each stage of development and the effective use of the software for its intended application. It is an asynchronous means of communication between all involved stakeholders, including end-users and future developers, which transforms the tacit knowledge of developers into an explicit, exchangeable format (Ding *et al.*, 2014; Rus and Lindvall, 2002).

We found that, in general, the documentation of the maps was insufficient to make explicit the tacit knowledge of the map developers. This presented a barrier to successfully and efficiently querying the SEMs assessed in our survey. We observed that mappers' knowledge of their data model, database structure and intended uses for their database were generally underreported in accompanying SEM study reports. Discussion dedicated to instructing end-users on how they could or should interact with the database was particularly limited. This might compromise the ability of nonspecialist users to query SEMs for their own research interests. Similarly, trends visualized and analyzed in SEM study reports, which might serve as illustrative examples of how to interact with the SEM, were not accompanied by any documentation of the queries used to obtain the analyzed subset of evidence from the database—apart from 1 instance where the authors referred to code in GitHub, but the link was broken (Cheng *et al.*, 2019).

A more common practice for facilitating end-user access to trends in the evidence base was the development of interactive data visualization dashboards (Bernes *et al.*, 2015). Unlike their underlying databases, these dashboards were generally accompanied by documentation detailing how users could interact with the dashboard. This interaction was intuitive and required minimal technical expertise—with many dashboards adopting “point-and-click” functionality. However, interactive visualization dashboards should not be conflated with the systematic map database itself. These dashboards represent the visualized

outputs of a set of predefined queries, where users can select which of the set to display. They can be thought of as user-interfaces which have been optimized for particular queries. However, users cannot devise and visualize customized queries through such dashboards. For this, access to the underlying database is required—reinforcing the need for its documentation.

Thus the role of high-quality software documentation in promoting transparency, growth, development and maintenance of SEMs as living evidence products should not be underestimated when adapting the methodology for EH.

Including database software capacity in evidence mapping teams. A final point of interest from this survey of current systematic mapping practice is that the multidimensionality of the relational database storage technology was not utilized in the CEE maps which employed the technology. This was evidenced by systematic maps which used a flat data structure even within a relational database software environment. Such maps included Neaves *et al.* (2015)—which presented a single, flat data table with expanded columns despite the authors' acknowledgment of the limitations of this structure and the capacity of the chosen storage technology to overcome them.

Reasons for implementing flat relational databases were unclear or unreported. However, facilitating the access of nonspecialist users to SEM outputs may have been a potential driver of this practice. Flat tables are associated with simple querying processes such as filtering columns, whereas relational databases require a more technically demanding process of constructing queries in structured query language (SQL). However, these concerns can readily be addressed by developing user-interfaces such as the visualization dashboards discussed above, and do not explain why inherently flat storage technologies, such as spreadsheets, were not used preferentially in such cases.

Thus, an alternative motivation for implementing flat relational databases might be a lack of familiarity with database storage technologies. This highlights a key challenge for adapting SEM methodology to EH, where subject specialists interested in mapping EH evidence may not have the necessary training to successfully implement graph-based storage. This underscores the value of comprehensive documentation—where the technical construction and querying of emerging maps might serve as training opportunities for others interested in the methodology. It also indicates the importance of developing these skills within mapping teams—where recruiting databasing specialists to SEM teams might be considered as important as recruiting statisticians to systematic review teams.

CONCLUSION

Systematic evidence mapping is an emerging methodology in EH. It offers a resource-efficient means of gaining valuable insights from a vast and rapidly growing evidence base. Its overarching aims, of organizing data and providing computational access to research, should facilitate evidence-based approaches to chemical risk assessment and risk management decision-making.

The methodology has been applied in the wider environmental sciences by the CEE. Characterizing the state-of-the-art of CEE systematic mapping practices offers valuable lessons for adapting the methodology for EH.

In particular, the rigid data structures which dominate current practice are ill-suited to the complex, heterogeneous and

highly connected data constituting the EH, and toxicology evidence bases. Flat data structures and those which are closely linked to predefined, on-write schema are optimized for a narrow range of specific use cases, which fits poorly with the much broader range of uses associated with chemicals policy workflows.

Successful adaptation of SEM methodology for EH would be accelerated by adopting flexible, schemaless database technologies in place of rigid, schema-first approaches. We have argued that knowledge graphs are 1 technological solution, which potentially provide an intuitive and scalable means of representing all of the connected, complex knowledge on a topic. Converse to the flat or relational databases favored by current practice, knowledge graphs store relationships between data as first-class entities, preserving their semantic value and making them accessible to queries. This ability to explore data through relationships or “patterns of information” does not require that users are familiar with a predefined data model or schema. This vastly expands the exploratory use cases of SEMs and even facilitates the discovery of new, previously uncharacterized relationships.

There are several readily available commercial and open-source graph database implementations (ArangoDB, 2019; Neo4j, 2019; Ontotext, 2019a; Stardog, 2019), and a variety of knowledge graph applications which demonstrate the power and utility of the graph data model and its inferencing capacity. Such resources are valuable for investigating the storage and exploration of SEMs as knowledge graphs and help to lower the entry barrier associated with familiarizing and training mappers in the use of a technology novel to the field.

SUPPLEMENTARY DATA

Supplementary data are available at *Toxicological Sciences* online.

ACKNOWLEDGMENTS

The authors would like to thank Mike Wolfe at Yordas Group for input helpful to the revision of this article.

FUNDING

T.A.M.W.'s PhD is financially supported by the Centre for Global Eco-innovation (a programme funded by the European Regional Development Fund) and Yordas Group, a global consultancy in the area of chemical safety, regulations, and sustainability. P.W.'s contribution to the manuscript was funded by the Evidence-Based Toxicology Collaboration at Johns Hopkins Bloomberg School of Public Health.

DECLARATION OF CONFLICTING INTERESTS

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

AUTHOR CONTRIBUTIONS

J.V. introduced the concept of graph databases, after which T.A.M.W., J.V., P.W., and C.H. established the principal ideas for the manuscript and developed an outline. T.A.M.W. wrote the first draft of the manuscript. T.A.M.W. and P.W.

conducted the survey of CEE systematic maps. J.V. offered technical expertise and edited the discussion accordingly. N.H. contributed to the revision of the manuscript, offering regulatory and chemical risk assessment expertise. All authors reviewed and edited the manuscript and contributed to its development.

REFERENCES

- Aggarwal, C. C., and Wang, H. (2010). *Managing and Mining Graph Data*, Vol. 40. Springer, New York. <https://doi.org/10.1007/978-1-4419-6045-0>.
- ArangoDB. (2019). *Graphs and ArangoDB*. Available at: <https://www.arangodb.com/arangodb-training-center/graphs/>. Accessed October 2019.
- Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., et al. (2000). The Gene Ontology Consortium, Michael Ashburner1, Catherine A. Ball3, Judith A. Blake4, David Botstein3, Heather Butler1, J. Michael Cherry3, Allan P. Davis4, Kara Dolinski3, Selina S. Dwight3, Janan T. Eppig4, Midori A. Harris3, David P. Hill4, Laurie Is. *Nat. Genet.* 25, 25–29.
- Barra Caracciolo, A., de Donato, G., Finizio, A., Grenni, P., Santoro, S., and Petrangeli, A. B. (2013). A new online database on chemicals in accordance with REACH regulation. *Hum. Ecol. Risk Assess.* 19, 1682–1699.
- Bernes, C., Bullock, J. M., Jakobsson, S., Rundlöf, M., Verheyen, K., and Lindborg, R. (2017). How are biodiversity and dispersal of species affected by the management of roadsides? A systematic map. *Environ. Evid.* 6, 1–16.
- Bernes, C., Jonsson, B. U., Junninen, K., Lohmus, A., Macdonald, E., Müller, J., and Sandström, J. (2015). What is the impact of active management on biodiversity in boreal and temperate forests set aside for conservation or restoration? A systematic map. *Environ. Evid.* 4, 1–22.
- Beverly, B. (2019). Abstract 3267: Potential Alternatives to Systematic Review: Evidence Maps and Scoping Reviews. Available at: <https://www.toxicology.org/events/am/AM2019/program-details.asp>. Accessed October 2019.
- Blaha, M. R., Premerlani, W. J., and Rumbaligh, J. E. (1988). Relational database design using an object-oriented methodology. *Comput. Pract.* 31, 414–427.
- Castellort, A., and Laurent, A. (2016). Rogue behavior detection in NoSQL graph databases. *J. Innov. Digit. Ecosyst.* 3, 70–82.
- Cheng, S. H., Macleod, K., Ahlroth, S., Onder, S., Perge, E., Shyamsundar, P., Rana, P., Garside, R., Kristjansson, P., McKinnon, M. C., et al. (2019). A systematic map of evidence on the contribution of forests to poverty alleviation. *Environ. Evid.* 8, 1–22.
- Clapton, J., Rutter, D., and Sharif, N. (2009). SCIE Systematic Mapping Guidance April 2009. Available at: <https://www.scie.org.uk/publications/researchresources/tr03.pdf>. Accessed October 2019.
- Collaboration for Environmental Evidence. (2019a). CEE Meetings. Available at: <https://www.environmentalevidence.org/meetings>. Accessed October 2019.
- Collaboration for Environmental Evidence. (2019b). Completed Reviews. Available at: <http://www.environmentalevidence.org/completed-reviews>. Accessed July 2019.
- Collaboration for Environmental Evidence. (2019c). *Environmental Evidence: Systematic Map Submission Guidelines*. Available at: <https://environmentalevidencejournal.biomedcentral.com/>

- submission-guidelines/preparing-your-manuscript/systematic-map. Accessed October 2019.
- Commission of the European Communities. (2001). *White Paper: Strategy for a Future Chemicals Policy*, Vol. 13. doi: 10.1007/BF03038641.
- Cresswell, C. J., Wilcox, A., Randall, N. P., and Cunningham, H. M. (2018). What specific plant traits support ecosystem services such as pollination, bio-control and water quality protection in temperate climates? A systematic map. *Environ. Evid.* 7, 1–13.
- Dillon, T., Chang, E., Hadzic, M., and Wongthongtham, P. (2008). Differentiating conceptual modelling from data modelling, knowledge modelling and ontology modelling and a notation for ontology modelling. In *Conferences in Research and Practice in Information Technology Series*, 79.
- Ding, W., Liang, P., Tang, A., and Van Vliet, H. (2014). Knowledge-based approaches in software documentation: A systematic literature review. *Inf. Softw. Technol.* 56, 545–567.
- ECHA. (2016). *Practical Guide How to Use Alternatives to Animal Testing to Fulfil Your Information Requirements for REACH Registration*.
- Edwards, S. W., Tan, Y.-M., Villeneuve, D. L., Meek, M. E., and McQueen, C. A. (2015). Adverse outcome pathways—Organizing toxicological information to improve decision making. *J. Pharmacol. Exp. Ther.* 356, 170–181.
- Elmasri, R., and Navathe, S. B. (2013). *The Relational Data Model and Relational Database Constraints. Fundamentals of Database Systems*. Pearson Education, UK.
- EPA. (2018). *Application of Systematic Review in TSCA Risk Evaluations*, pp. 1–247. Available at: https://www.epa.gov/sites/production/files/2018-06/documents/final_application_of_sr_in_tsca_05-31-18.pdf.
- European Food Safety Authority. (2010). Application of systematic review methodology to food and feed safety assessments to support decision making¹. EFSA Guidance for those carrying out systematic reviews. *EFSA J.* 8, 1637.
- Ghrab, A., Romero, O., Skhiri, S., Vaisman, A., and Zimányi, E. (2016). GRAD: On Graph Database Modeling. Available at: <http://arxiv.org/abs/1602.00503>. Accessed October 2019.
- Gupta, P., Goel, A., Lin, J., Sharma, A., Wang, D., and Zadeh, R. (2013). WTF: The Who to Follow Service at Twitter, WWW '13: Proceedings of the 22nd international conference on World Wide Web, pp. 505–514. 10.1145/2488338.2488433.
- Haddaway, N. R., Bernes, C., Jonsson, B.-G., and Hedlund, K. (2016). The benefits of systematic mapping to evidence-based environmental management. *AMBIO* 45, 613–620.
- Haddaway, N. R., Brown, C., Eales, J., Eggers, S., Josefsson, J., Kronvang, B., Randall, N. P., and Uusi-Kämppe, J. (2018a). The multifunctional roles of vegetated strips around and within agricultural fields. *Environ. Evid.* 7, 1–43.
- Haddaway, N. R., and Crowe, S. (eds) (2018). *Stakeholder Engagement in Environmental Evidence Synthesis*. Available at: <http://www.eviem.se/Documents/projekt/2018/SRbookAll.pdf>. Accessed October 2019.
- Haddaway, N. R., Hedlund, K., Jackson, L. E., Kätterer, T., Lugato, E., Thomsen, I. K., Jørgensen, H. B., and Söderström, B. (2015). What are the effects of agricultural management on soil organic carbon in boreo-temperate systems? *Environ. Evid.* 4, 1–29.
- Haddaway, N. R., Macura, B., Whaley, P., and Pullin, A. S. (2018b). ROSES Reporting standards for Systematic Evidence Syntheses: Pro forma, flow—Diagram and descriptive summary of the plan and conduct of environmental systematic reviews and systematic maps. *Environ. Evid.* 7, 4–11.
- Haddaway, N. R., Styles, D., and Pullin, A. S. (2014). Evidence on the environmental impacts of farm land abandonment in high altitude/mountain regions: A systematic map. *Environ. Evid.* 3, 17–19.
- Hardy, B., Apic, G., Carthew, P., Clark, D., Cook, D., Dix, I., Escher, S., Hastings, J., Heard, D. J., Jeliakova, N., et al. (2012). Toxicology ontology perspectives. *ALTEX* 29, 139–156.
- Have, C.T., and Jensen, L. J. (2013). Databases and ontologies are graph databases ready for bioinformatics? *Bioinformatics* 29, 3107–3108.
- Weaver, J and Tarjan, P. (2013). Facebook linked data via the Graph API. *Semant. Web.* 4, 245–250.
- Hoffmann, S., de Vries, R. B. M., Stephens, M. L., Beck, N. B., Dirven, H. A. A. M., Fowle, J. R., Goodman, J. E., Hartung, T., Kimber, I., Lalu, M. M., et al. (2017). A primer on systematic reviews in toxicology. *Arch. Toxicol.* 91, 2551–2575.
- Hoffmann, S., and Hartung, T. (2006). Toward an evidence-based toxicology. *Hum. Exp. Toxicol.* 25, 497–513.
- Ingre-Khans, E., Ågerstrand, M., Beronius, A., and Rudén, C. (2016). Transparency of chemical risk assessment data under REACH. *Environ. Sci.: Process. Impacts* 18, 1508–1518.
- James, K. L., Randall, N. P., and Haddaway, N. R. (2016). A methodology for systematic mapping in environmental sciences. *Environ. Evid.* 5, 7.
- Janković, S., Mladenović, S., Mladenović, D., Vesković, S., and Glavić, D. (2018). Schema on read modeling approach as a basis of big data analytics integration in EIS. *Enterp. Inf. Syst.* 12, 1180–1201.
- Johnson, V., Fitzpatrick, I., Floyd, R., Simms, A. (2011). What is the evidence that scarcity and shocks in freshwater resources cause conflict instead of promoting collaboration? CEE review 10-010. Collaboration for Environmental Evidence, Bangor, UK.
- Karp, P. D. (1996). Database links are a foundation for interoperability. *Trends Biotechnol.* 14, 273–279.
- Kleppmann, M. (2017). *Designing Data-intensive Applications*. O'Reilly Media, Inc, UK. Available at: <https://www.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/>.
- Leisher, C., Tamsah, G., Booker, F., Day, M., Samberg, L., Prosnitz, D., Agarwal, B., Matthews, E., Roe, D., Russell, D., et al. (2016). Does the gender composition of forest and fishery management groups affect resource governance and conservation outcomes? A systematic map. *Environ. Evid.* 5, 1–10.
- Lewis, K. A., Tzilivakis, J., Warner, D. J., and Green, A. (2016). An international database for pesticide risk assessments and management. *Hum. Ecol. Risk Assess.* 22, 1050–1064.
- Liu, Z. H., and Gawlick, D. (2015). *Management of Flexible Schema Data in RDBMSs—Opportunities and Limitations for NoSQL in CIDR 2015*.
- Lyndon, M. (1989). Information economics and chemical toxicity: Designing laws to produce and use data. *Mich. Law Rev.* 87, 1795–1861.
- Macura, B., Secco, L., and Pullin, A. S. (2015). What evidence exists on the impact of governance type on the conservation effectiveness of forest protected areas? Knowledge base and evidence gaps. *Environ. Evid.* 4, 24.
- Mandrioli, D., Schlünssen, V., Ádám, B., Cohen, R. A., Colosio, C., Chen, W., Fischer, A., Godderis, L., Göen, T., Ivanov, I. D., et al. (2018). WHO/ILO work-related burden of disease and injury: Protocol for systematic reviews of occupational exposure to dusts and/or fibres and of the effect of occupational exposure

- to dusts and/or fibres on pneumoconiosis. *Environ. Int.* **119**, 174–185.
- Martin, O. V., Geueke, B., Groh, K. J., Chevrier, J., Fini, J.-B., Houlihan, J., Kassotis, C., Myers, P., Nagel, S. C., Pelch, K. E., et al. (2018). Protocol for a systematic map of the evidence of migrating and extractable chemicals from food contact articles. 10.5281/zenodo.2525277.
- Martin, P., Bladier, C., Meek, B., Bruyere, O., Feinblatt, E., Touvier, M., Watier, L., and Makowski, D. (2018). Weight of evidence for hazard identification: A critical review of the literature. *Environ. Health Perspect.* **126**, 076001–076015.
- Mcintosh, E. J., Chapman, S., Kearney, S. G., Williams, B., Althor, G., Thorn, J. P. R., Pressey, R. L., McKinnon, M. C., and Grenyer, R. (2018). Absence of evidence for the conservation outcomes of systematicconservation planning around the globe: a systematic map. *Environ. Evid.* **7**, 22.
- McKinnon, M. C., Cheng, S. H., Dupre, S., Edmond, J., Garside, R., Glew, L., Holland, M. B., Levine, E., Masuda, Y. J., Miller, D. C., et al. (2016). What are the effects of nature conservation on human well-being? A systematic map of empirical evidence from developing countries. *Environ. Evid.* **5**, 1–25.
- National Center for Biomedical Ontology. (2019). *BioPortal*. Available at: <https://bioportal.bioontology.org/>. Accessed October 2019.
- Neaves, L. E., Eales, J., Whitlock, R., Hollingsworth, P. M., Burke, T., and Pullin, A. S. (2015). The fitness consequences of inbreeding in natural populations and their implications for species conservation—A systematic map. *Environ. Evid.* **4**, 1–17.
- Neo4j. (2019). *Neo4j*. Available at: <https://neo4j.com/>. Accessed October 2019.
- Noy, N. F., and Klein, M. (2004). Ontology evolution: Not the same as schema evolution. *Knowl. Inf. Syst.* **6**, 428–440.
- NTP-OHAT. (2019). *About the Office of Health Assessment and Translation*. Available at: <http://ntp.niehs.nih.gov/>. Accessed October 2019.
- Oliver, S., and Dickson, K. (2016). Policy-relevant systematic reviews to strengthen health systems: Models and mechanisms to support their production. *Evid. Policy* **12**, 235–259.
- Ontotext. (2019a). *Ontotext GraphDB*. Available at: <https://www.ontotext.com/products/graphdb/>. Accessed October 2019.
- Ontotext. (2019b). *What is a Knowledge Graph?* Available at: <https://www.ontotext.com/knowledgehub/fundamentals/what-is-a-knowledge-graph/>. Accessed October 2019.
- Ontotext. (2019c). *What is Inference?* Available at: <https://www.ontotext.com/knowledgehub/fundamentals/what-is-inference/>. Accessed October 2019.
- Pavlopoulos, G. A., Secier, M., Moschopoulos, C. N., Soldatos, T. G., Kossida, S., Aerts, J., Aerts, J., Schneider, R., and Bagos, P. G. (2011). Using graph theory to analyze biological networks. *BioData Min.* **4**, 1–27.
- Pelch, K. E., Bolden, A. L., and Kwiatkowski, C. F. (2019). Environmental chemicals and autism: A scoping review of the human and animal research. *Environ. Health Perspect.* **127**, 046001.
- Pelch, K. E., Reade, A., Wolffe, T. A. M., and Kwiatkowski, C. F. (2019). PFAS health effects database: Protocol for a systematic evidence map. *Environ. Int.* **130**, 104851.
- Pool, R., and Rusch, E. (2014). *Identifying and Reducing Environmental Health Risks of Chemicals in Our Society: Workshop Summary*. The National Academies Press, Washington, DC.
- Randall, N. P., Donnison, L. M., Lewis, P. J., and James, K. L. (2015). How effective are on-farm mitigation measures for delivering an improved water environment? A systematic map. *Environ. Evid.* **4**, 18.
- Randall, N. P., and James, K. L. (2012). The effectiveness of integrated farm management, organic farming and agri-environment schemes for conserving biodiversity in temperate Europe—A systematic map. *Environ. Evid.* **1**, 4.
- Rhomberg, L. R., Goodman, J. E., Bailey, L. A., Prueitt, R. L., Beck, N. B., Bevan, C., Honeycutt, M., Kaminski, N. E., Paoli, G., Pottenger, L. H., et al. (2013). A survey of frameworks for best practices in weight-of-evidence analyses. *Crit. Rev. Toxicol.* **43**, 753–784.
- Robinson, I., Webber, J., and Eifrem, E. (2015). *Graph Databases*. O'Reilly Media, New York. doi: 10.1016/B978-0-12-407192-6.00003-0.
- Rus, I., and Lindvall, M. (2002). Knowledge management in software engineering. *IEEE Softw.* **19**, 26–38.
- Sadowski, G., and Rathle, P. (2015). *Fraud Detection: Discovering Connections with Graph Databases*. Neo Tehcnology. Available at: <https://neo4j.com/use-cases/fraud-detection/>. Accessed October 2019.
- Schultz, T. W., Cronin, M. T. D., Walker, J. D., and Aptula, A. O. (2003). Quantitative structure-activity relationships (QSARS) in toxicology: A historical perspective. *J. Mol. Struct.: THEOCHEM* **622**, 1–22.
- Segaran, T., Evans, C., and Taylor, J. (2009). *Programming the Semantic Web: Traditional Data-modeling Methods*. O'Reilly Media, USA.
- Sexton, K., and Hattis, D. (2007). Assessing cumulative health risks from exposure to environmental mixtures—Three fundamental questions. *Environ. Health Perspect.* **115**, 825–832.
- Singhal, A. (2012). *Introducing the Knowledge Graph: Things, Not Strings*. Available at: <https://googleblog.blogspot.com/2012/05/introducing-knowledge-graph-things-not.html>. Accessed October 2019.
- Sola, P., Cerutti, P. O., Zhou, W., Gautier, D., Iiyama, M., Shure, J., Chenevoy, A., Yila, J., Dufe, V., Nasi, R., et al. (2017). The environmental, socioeconomic, and health impacts of woodfuel value chains in Sub-Saharan Africa: A systematic map. *Environ. Evid.* **6**, 1–16.
- Stardog. (2019). *Stardog*. Available at: <https://www.stardog.com/>. Accessed October 2019.
- The Endocrine Disruption Exchange. (2019). *TEDX Publications*. Available at: <https://endocrinedisruption.org/interactive-tools/publications/>. Accessed October 2019.
- Thorn, J. P. R., Friedman, R., Benz, D., Willis, K. J., and Petrokofsky, G. (2016). What evidence exists for the effectiveness of on-farm conservation land management strategies for preserving ecosystem services in developing countries? A systematic map. *Environ. Evid.* **5**, 1–29.
- United States Environmental Protection Agency. (2016). *The Frank R. Lautenberg Chemical Safety for the 21st Century Act*. Available at: <https://www.epa.gov/assessing-and-managing-chemicals-under-tsca/frank-r-lautenberg-chemical-safety-21st-century-act/>. Accessed October 2019.
- Vandenberg, L. N., Ågerstrand, M., Beronius, A., Beausoleil, C., Bergman, Å., Bero, L. A., Bornehag, C.-G., Boyer, C. S., Cooper, G. S., Cotgreave, I., et al. (2016). A proposed framework for the systematic review and integrated assessment (SYRINA) of endocrine disrupting chemicals. *Environ. Health* **15**, 1–19.
- Vinken, M., Whelan, M., and Rogiers, V. (2014). Adverse outcome pathways: Hype or hope? *Arch. Toxicol.* **88**, 1–2.
- Walker, V. R., Boyles, A. L., Pelch, K. E., Holmgren, S. D., Shapiro, A. J., Blystone, C. R., Devito, M. J., Newbold, R. R., Blain, R., Hartman, P., et al. (2018). Human and animal evidence of potential transgenerational inheritance of health effects: An evidence map and state-of-the-science evaluation. *Environ. Int.* **115**, 48–69.

- Webber, J. (2018). *Powering Real-time Recommendations With Graph Database Technology Powering Real-time*. Neo4j.
- Wolffe, T. A. M., Whaley, P., Halsall, C., Rooney, A. A., and Walker, V. R. (2019). Systematic evidence maps as a novel tool to support evidence-based decision-making in chemicals policy and risk management. *Environ. Int.* **130**, 104871.
- World Health Organization. (2019). *Framework for Use of Systematic Review Methods in Chemical Risk Assessment—Authors Meeting*. Available at: http://who.int/ipcs/events/SRmeeting_US/en/. Accessed October 2019.
- Zynda, M. (2013). The first killer app: A history of spreadsheets. *Interactions* **20**, 68–72.